

Law, Social Welfare, and Net Neutrality

Keith N. Hylton^{1,2}

© Springer Science+Business Media New York 2016

Abstract Net neutrality generates wealth transfers from one type of internet content provider to another. In theory, these transfers might be socially desirable, and could be justified on the basis of informational externalities similar to those cited to justify fair use in copyright law. In practice, however, the conditions that justify fair use do not hold where net neutrality operates. Moreover, the internal subsidization required by net neutrality generates a regressive transfer. The welfare gains that might come from controlling anticompetitive abuse or government coercion through implementation of net neutrality can be achieved by alternative policies with less harmful consequences.

Keywords Anticompetitive exclusion · Fair use · Net neutrality · Telecommunications regulation · Vertical integration

Net neutrality is an appealing term. After all, what might be the conceptual antithesis of net neutrality: net discrimination? The word discrimination suggests an evil purpose, while neutrality suggests immediately an open, fair process. The term alone conditions one to support the concept.

Net neutrality consists, as far as I can tell, of a regulatory norm that is as simple as it is appealing. Providers of content transported through the networks of broadband internet firms cannot be charged different prices by the broadband firms for the service of transporting their content.¹ Thus, a firm that provides highly

¹ As Becker et al. (2010, p. 498) explain, the net neutrality principle has devolved into four specific requirements: broadband providers are “prohibited from: (1) prioritizing traffic and charging differential

✉ Keith N. Hylton
knhylton@bu.edu

¹ William Fairfield Warren Distinguished Professor, Boston University, Boston, MA, USA

² Boston University School of Law, 765 Commonwealth Avenue, Boston, MA 02215, USA

demand content that absorbs much of the capacity of the broadband network cannot be charged for the additional congestion and wear-and-tear that is associated with the transporting of its content.

If there were no differences between the consumers of various services—for example, if everyone consumed the same information services from the internet—net neutrality would truly be neutral in effect. It would not permit differential pricing, or discrimination, to adversely affect any providers of content. It would not differentially impact any consumer, since every consumer is identical by hypothesis.

But consumers of internet services are not identical: they demand different services, which impose dissimilar costs on broadband firms.² Consequently, net neutrality requires some consumers to subsidize the consumption of others. In this sense, net neutrality is not neutral at all: it forces A to pay for the consumption of B. Viewed from this perspective, net neutrality is a form of differential pricing.

The basic economics of net neutrality seems similar to that of toll bridges.³ Charging all users of the bridge the same amount may force some users to subsidize others. This has moral hazard and adverse selection implications. Consumer welfare declines, unless there is an efficiency case for a regulation that requires internal subsidization of one group of consumers by another group of consumers (cross-subsidization).⁴ I will consider the possible welfare cases for such internal subsidization here, drawing on familiar analogies in the law. After that, I will consider alternative arguments for the net neutrality norm based on distributional concerns, potential anticompetitive abuse, and government coercion.

I conclude that the case for net neutrality is weak: the potential efficiency justifications for the policy are speculative, and unsupported so far by the evidence. The distributional consequences of the policy are undesirable. The policy effects a transfer of resources from the less advantaged to the more advantaged without any significant offsetting welfare gain. Every speculative gain that might come from the policy—such as controlling anticompetitive abuse or government coercion—can be achieved by an alternative policy with less harmful consequences.

Footnote 1 continued

prices based on the priority status; (2) imposing congestion-related charges; (3) adopting business models that offer exclusive content or that establish exclusive relationships with particular content providers; and (4) charging content providers to access the Internet based on factors other than the bandwidth supplied”.

² On the variation in network services and costs, see Yoo (2013).

³ The similarity is that in both contexts, one observes large up-front capital expenses to construct a facility, and the incremental cost of using the facility is comparatively small. The economics of bridges is a topic with an extensive literature. See e.g., Hotelling (1938) and Minasian (1979). The difficult problem of optimal supply that was analyzed in the earlier literature has important implications for the treatment of different types of users, and the extent to which one can determine when the pricing of bridge services inefficiently advantages one user at the expense of another user. I will oversimplify and avoid most of these difficulties here.

⁴ I use the term cross-subsidy as an equivalent to “internal subsidy”, as in Posner (1971). This is a popular, though admittedly non-rigorous use of the term. In my treatment, a cross-subsidy or internal subsidy results from a regulation (such as a net neutrality requirement) when a firm is forced to change its pricing structure, as a result of the regulation, in a manner that benefits one group of consumers and disadvantages another group of consumers. The concept of cross-subsidy is difficult and would require a much more extensive treatment than is offered here to deal with adequately in economic terms. A rigorous treatment of the concept of cross-subsidy is provided in Faulhaber (1975).

1 The Bridge Analogy

Consider a toll bridge with two types of users: ordinary cars, and heavy trucks. The difference between the two is that the heavy trucks are more costly to the bridge owner to service because they impose more congestion on bridge traffic. Congestion reduces the flow of traffic over the bridge, thereby increasing the average cost of the service. In addition, the trucks impose more depreciation (wear and tear) on the bridge. A profit-maximizing bridge owner would adopt a system of “discriminatory” Ramsey prices: the owner would charge a mark-up above user-specific marginal cost that is inversely related to the elasticity of user demand.⁵ If we assume that the demand elasticity of cars is no greater than that of trucks—cars involve a larger percentage of leisure travelers who are more likely to seek cheaper though more time-consuming routes—the owner would charge higher prices to trucks than to cars.

A regulatory authority that seeks to maximize welfare subject to a given level of profit that is promised to the bridge owner would also choose a system of discriminatory Ramsey prices, though not precisely the same as those chosen by the unregulated monopolist. The mark-up above user-specific marginal cost would be somewhat less, but would still vary inversely with the elasticity of user demand.⁶

Suppose, however, that the regulator requires the bridge owner to charge the same price to both cars and trucks. Under the conditions assumed, cars would, in effect, finance an internal subsidy for trucks. As a result, trucks would tend to use the bridge more than if the cross-subsidy were not present, and cars would use the bridge less than if the cross-subsidy were not present. The additional use by trucks would increase congestion costs and thereby increase the cost of the service and increase wear-and-tear costs.

Charging cars and trucks different prices would permit the bridge owner to internalize to truck owners the additional costs that are imposed by the trucks. This, in turn, would discourage the trucks from excessive use—for example, from imposing a marginal cost of \$1 on the bridge owner and other users when the marginal benefit to the truck owner from the particular use is only \$0.50. A charge that varied with the intensity of the use would encourage truck owners to consider the congestion costs and the miles of wear and tear that are imposed in each relevant time period. The higher charge would also induce some truck owners to avoid the bridge in favor of another route. Over time, charges might encourage technological innovation toward trucks that carry the same freight while imposing lower congestion and depreciation costs.

Charging separate prices allows the bridge owner to reduce congestion and depreciation costs, and pass those cost savings on to consumers in the form of lower general prices (for an equivalent unit of service) for use of the bridge, which, in turn, would increase the total consumption of the services that are offered by the bridge.

⁵ Here I refer to the simplest form of Ramsey pricing that results from profit maximization by a multi-product monopolist. For simplicity, I assume that the user demand functions are independent of each other. For a straightforward presentation, see Tirole (1980).

⁶ See generally Baumol and Bradford (1970).

Admittedly, in some cases the bridge owner might choose not to charge differential prices. Perhaps the differences in service costs are minor, and the administrative costs of differential pricing exceed the efficiency gains. Alternatively, perhaps trucks provide the greatest source of demand for new bridge capacity. Foresighted bridge owners would therefore be reluctant to tax a major source of industrial capacity growth. In these cases, the bridge owner may choose not to impose differential pricing even if completely free to do so.

The bridge analogy seems to apply straightforwardly to the net neutrality problem: net neutrality is equivalent to prohibiting the bridge owner from using differential pricing, and generates similar costs. Some broadband-intensive providers of internet content, such as Netflix, would—in the absence of differential pricing—impose extraordinary congestion costs that would result in an internal subsidy from consumers of other internet services.⁷ Hence, permitting the network owner to price differentially can and probably would enhance consumer welfare.⁸ To the extent that heavy use of the service has a depreciation effect (electrical components suffer wear and tear from use), similar costs are imposed.⁹

2 Is There a Case for Cross-Subsidization?

As a general matter, it is not difficult to justify cross-subsidization in the presence of significant externalities; but the question is whether such a justification is applicable in the net neutrality context.

Return to the bridge story: if trucks provide a positive externality to all bridge users, then cross-subsidization might be socially optimal. Suppose, for example, the trucks provide a good that benefits other bridge users (or the consumers of those other bridge users) through a mechanism that fails to charge a price for those external benefits. Vaccines, for example, provide a beneficial externality to consumers who do not consume the vaccines. A subsidy toward the trucks that carry those vaccines might be optimal. If most of the trucks carry vaccines, a subsidy that is funded by cars and that benefits trucks in general might be optimal.

Similarly, if the services that impose the greatest congestion costs on the internet also provide a positive externality to all consumers of internet services, then the cross-subsidies that are created by net neutrality might be socially optimal. The informational services that are often cited for providing important positive externalities to society are educational. Educational services enhance the stock of human capital, and thereby help to make members of society more productive, or otherwise advance science and culture. Because of these benefits, the law has in some cases provided an implicit subsidy for informational services that are primarily educational.

⁷ See Clark (2014).

⁸ See e.g., Becker et al. (2010) and Peitz and Schuett (2015). In my example in the text, if the privately-chosen prices are relatively close to the socially optimal Ramsey prices, then unregulated monopoly would enhance welfare relative to the equal-pricing rule.

⁹ Although the empirical importance of such depreciation may be entirely speculative, for a model that incorporates depreciation as a factor, see Odlyzko (1997).

The fair use doctrine of copyright law, for example, imposes a subsidy that is funded by one type of information service for the benefit of another type of information service: from the holder of a copyright to an infringer of the copyright. Under the fair use doctrine, an agent who copies copyrighted material is not infringing if the use is deemed fair. Fair use is determined by several factors, such as the extent and purpose of the copying.

One key factor that supports a finding of fairness is copying for educational purposes. Educational purposes can be understood broadly to include information that advances culture. Satire, for example, arguably has an educational purpose, because it is through satire that the public becomes aware of the shortcomings or flaws of certain theories or works of art.

For example, Alice Randall's *The Wind Done Gone* provides an account, deemed a parody by one court, of the fictional events from Margaret Mitchell's *Gone with the Wind*, told from the eyes of slaves rather than from the eyes of slaveholders. In a society that has matured to the stage that its prior belief that one particular race should be given the legal power to hold another race in slavery now seems comically unjust, Randall's retelling, from an upside-down perspective, of a famous work of art embodying the now-discredited view of social order has significant educational value. By subsidizing such iconoclastic works of literature, fair use doctrine probably enhances society's welfare.¹⁰

Cross-subsidization of the sort that is required under the net neutrality principle is therefore not foreign to the common law, which for the most part strongly protects property rights, nor necessarily harmful to social welfare. The question is whether this fair-use based theory of cross-subsidization helps justify arguments that support net neutrality.

If the providers of heavily demanded internet content were providing educational services, or informational services that generally enhance the stock of human capital, then the net neutrality norm might lead to a socially optimal regulatory regime. However, as far as I am aware, this is not the case.

Netflix, for example, streams popular movies and television shows through the internet, using an extraordinary amount of pipe capacity as it does so. Some of the shows are critically acclaimed, but I doubt that any of them could be accurately described as educational in any significant sense.

The most popular Netflix program, *House of Cards*, offers a realistically cynical view of the inner workings of American politics, but such depictions of the political system have been part of western literature for a long time and probably do little to enhance the stock of human capital significantly. To be sure, Shakespeare also gave us realistically cynical portraits of the inner workings of politics, but the similarity between Shakespeare and *House of Cards* is probably equivalent to the similarity between a Rembrandt and a typical piece of critically-acclaimed modern art: one truly educates and significantly enhances culture, while the other mainly reflects the current tastes of the median consumer.

Economides and Tåg (2012) offer the most prominent externality justification for network neutrality to date, though it is different from the fair-use-based theory that

¹⁰ Gordon (1982).

was just described. Instead of externalities across different types of consumer, the external effect may be network externalities between consumer and content provider. Network effects may work in both directions across the platform, with more consumers generating more content, and more content generating more consumers. The broadband network provider may tend to overcharge the content provider relative to the social optimum. If the network externality that is suggested by Economides and Tag is sufficiently strong, implementation of the network neutrality norm might result in an optimal outcome.

Although the Economides and Tag result is rather weak—under special conditions net neutrality might be socially optimal—there are still reasons to question its applicability to the real world. The result requires strong assumptions, such as identical consumers and inelastic demand. However, much of the economic controversy that surrounds network neutrality stems from the heterogeneity of consumers and the existence of some degree of demand elasticity for network services. The bridge analogy with which this paper begins assumes these features. Moreover, the empirical evidence with regard to the importance of provider-to-consumer network effects is speculative at this stage, and the theory may not capture all of the relevant externalities (Becker et al. 2010).

3 Distributional Considerations

On distributional grounds, cross-subsidizing informational services can harm social welfare, at least under a Rawlsian perspective,¹¹ which disfavors wealth transfers from the materially advantaged to the relatively disadvantaged. A Rawls-inspired welfare function would incorporate the distribution of resources as a component of social welfare.¹² A more equal distribution of resources reduces the frequency of instances where individuals are born into, or somehow find themselves through no fault of their own, in impoverished households. Thus, institutions that reduce the risk of being placed into such households enhance society's welfare if members of society are sufficiently risk averse, and conversely institutions that increase this risk reduce society's welfare. This perspective suggests efficiency can be decomposed into general efficiency (or "wealth maximization") and distributional efficiency concerns. I will focus on the distributional efficiency concern here.

On the distributional effects of imposing cross-subsidies in information services, consider the British Broadcasting Corporation (BBC). Because of the greater educational value of its programming, the BBC generally has a stronger claim to subsidies than does Netflix, yet the BBC subsidy is regressive in distributional terms. It reduces the price of a service that is consumed mainly by the wealthy and compels the poor, through its government funding, to pay part of the cost of providing those services to the rich.

¹¹ Rawls (1971).

¹² See e.g., Posner (1981, p. 59). Of course, one could say that this is just an ordinary social welfare function that takes risk-aversion into account. See id.

Similarly, to the extent that net neutrality results in subsidizing Netflix, it imposes a regressive tax on society. The Netflix programs are less high-brow than those of the BBC, and as a result have a greater appeal to the average middle-income consumer, but the poor are less likely to be among the heavy consumers of Netflix.

A recent study of the demographics of Netflix users notes that “when compared to the average U.S. adult, recent Netflix subscribers are:

- (a) Gender: fairly even gender split, 49% male, 51% female;
- (b) Age: 37% are Millennials (18–34), which is 23% more likely than average;
- (c) Residential status: slightly more likely to be adults who still live with their parents, although the majority (61%) of recent subscribers are homeowners;
- (d) Suburban living: 50% live in the suburbs, which is slightly more than average (+9%);
- (e) Household income (HHI): 42% have a HHI of \$50K and under, 35% have a HHI of \$50K–\$100K, and 23% have a HHI over \$100 K;
- (f) Education: they are slightly more likely than average (+11%) to say their highest level of education is an associate or bachelor’s degree;
- (g) Parental status: 24% more likely to be parents and 52% more likely to have school-aged children living with them; and
- (h) Household: 38% more likely to have 4 or more people living in their current household.”¹³

These factors suggest that the average Netflix consumers is wealthier than the average U.S. citizen. On distributional grounds, net neutrality therefore would operate as a regressive tax by compelling the materially less-advantaged to support the consumption of a relatively advantaged group.

As a general rule, a norm that prohibits forced wealth transfers from the relatively disadvantaged to the relatively advantaged appears to have a stronger appeal on welfare grounds than the net neutrality norm. The policy against regressive transfers is likely to improve social welfare, when the distribution of wealth is taken into account, while the net neutrality norm degrades distributional efficiency. If net neutrality enhanced general efficiency—that is, enhanced society’s wealth—then it might be desirable in spite of its negative distributional effects. But the general efficiency argument for net neutrality is ambiguous, lacking empirical support currently, and unlikely to be proven valid even when the empirical evidence has been fully explored. Hence, the net neutrality norm at present offers an unambiguous reduction in distributional efficiency that is coupled with a likely negative impact on general efficiency.

The norm against regressive coerced transfers operates generally, while the net neutrality norm operates within a narrow subset of society. This suggests an important difference between the two norms, with distributional implications: the policy that disfavors regressive transfers is unlikely to serve effectively as a veil for concealing the aims of some faction that is interested in transferring wealth to itself,

¹³ <http://civicscience.com/wp-content/uploads/2015/05/CivicScience-Netfixs-New-Users-May-2015-Final.pdf>.

because that policy operates across a wide spectrum of society. A narrow faction that appeals to the anti-regressiveness principle purely out of self-interest—that is, to justify efforts to expropriate wealth from others—could easily find itself losing as a result of the principle just as much as it gains.

The net neutrality norm, by contrast, offers a narrowly tailored argument that, by seeming to invoke a lofty goal, obscures self-interested motives. The beneficiaries of net neutrality form a concentrated interest group, and therefore would benefit from a norm that conceals an effort to persuade government to transfer wealth in its favor. Given the net neutrality norm's usefulness as an instrument that simultaneously facilitates and obscures expropriative activity in the political process, it should be regarded with some suspicion from a distributional efficiency perspective.

4 Vertical Integration

One reason that the bridge analogy may seem incomplete as stated previously is that in some settings the ability to price differentially among bridge users might be used in an anticompetitive manner. Suppose one truck service purchases the bridge and charges high prices to rival truckers to use the bridge. This scenario is often described as the essential facilities problem in antitrust, where the classic case, *United States v. Terminal Railroad*,¹⁴ involved a consortium of railroads that purchased a bridge (actually, the rail terminal facilities that were connected to the bridge) that had to be used by competing railroads. In the internet network setting, a vertically integrated network owner (integrated into content) might use differential pricing to harm competitors in the content market.¹⁵

I have already established that differential pricing could encourage efficient (in a second-best sense) consumption of bridge services, and the same is true of general platforms such as a broadband network. Differential pricing might also provide the most efficient method of recouping the costs of constructing a platform, such as a bridge or a broadband network.¹⁶ The net neutrality norm is therefore inadvisable because of its efficiency costs. However, net neutrality might be defended still, because it prevents anticompetitive price discrimination among platform users.

There are reasons to question this argument for net neutrality: first, net neutrality goes further than necessary. Antitrust laws already exist for regulating anticompetitive conduct, and they attempt to regulate with a finer brush than the net neutrality rule. An antitrust court would take efficiencies into account in any analysis of a complaint against a vertically-integrated platform owner on antitrust grounds. The net neutrality principle ignores efficiencies.

Second, the market provides disincentives to predatory abuse in the vertical integration setting. If vertical integration yields no efficiencies, the vertically-integrated bridge owner will be punished, to some degree, by capital markets as investors observe the risk that the vertically integrated owner might reduce the

¹⁴ 224 U.S. 383 (1912).

¹⁵ See e.g., Hazlett and Wright (2012).

¹⁶ On Ramsey pricing and broadband networks, see Yoo (2013, p. 580).

market value of the bridge in order to provide an advantage to its own content flowing over the bridge.

This capital-markets reaction may seem speculative, at first, but examples exist: Keurig, a maker of coffee-brewing devices that is famously associated with the small plastic cups (K-cups) that are used for making coffee, recently abandoned its effort to lock out the products of rival K-cup makers from compatibility with its brewing device. When Keurig attempted to switch to a new brewing device that made rival K-cups technologically incompatible, the firm's stock price dropped so quickly that it reversed course and announced that it would continue marketing a device that could accommodate rival cups. Moreover, this reversal happened despite the fact that Keurig had virtually no significant competition in the market for its brewing devices.

The Keurig experience suggests that a firm that owns a facility that rivals must access to provide service to customers may suffer a severe penalty in the capital markets for attempting to exclude rival service providers. The value of the facility is determined by the overall quality and variety of services that the facility can provide. Locking out rival service providers may reduce the value of the facility. Hence, predatory exclusion is not always a costless exercise, even in the absence of the threat of competition or antitrust penalties; and in the case of Keurig this exclusion would have been permitted by the capital markets to occur only if it were efficient.

Broadband firms face the same capital market constraints as the makers of coffee brewing devices. If content exclusion were an efficient strategy for broadband firms, capital markets would steer them toward it. But remarkably few instances that potentially could be described as anticompetitive content exclusion have been observed.¹⁷

Finally, direct competition with other platform providers constrains anticompetitive abuse. While high-speed broadband networks tend to face few competitors in most American markets,¹⁸ the existence of geographic markets in which some consumers have access to more than one high-speed network suggests that the threat of competitive entry exists in the broadband market. Google's fiber service exists in eight metropolitan areas,¹⁹ and after initially promising to expand to more areas,²⁰ Google has changed course to pursue alternative evolving technologies.²¹

¹⁷ See Hazlett and Wright (2012, p. 782).

¹⁸ Most people have only one high-speed broadband provider; see <http://consumerist.com/2014/12/18/govt-report-true-high-speed-broadband-competition-in-the-u-s-remains-largely-nonexistent/>.

¹⁹ See Google Curbs Expansion of Fiber Optic Network, Cutting Jobs, New York Times, http://www.nytimes.com/2016/10/26/technology/google-curbs-expansion-of-fiber-optic-network-cutting-jobs.html?_r=0.

²⁰ See <http://consumerist.com/2015/10/29/shortlist-of-future-google-fiber-cities-keeps-getting-longer-three-more-cities-added/>.

²¹ See Google Curbs Expansion of Fiber Optic Network, Cutting Jobs, New York Times, http://www.nytimes.com/2016/10/26/technology/google-curbs-expansion-of-fiber-optic-network-cutting-jobs.html?_r=0. Another new alternative is the use of small satellites to provide internet service: <http://consumerist.com/2015/06/11/the-guy-who-started-tesla-wants-to-shoot-the-internet-to-you-from-space/>.

Moreover, the most heavily demanded content providers, which create the greatest congestion costs, also drive a substantial share of the demand for high-speed networks; this is an external effect that broadband providers have strong incentives to consider. An anticompetitive move to disadvantage such content, merely to provide a short-term boost in content that is owned by the vertically-integrated network, would risk a larger loss in market share in the future as competitive substitutes develop.

The recently-approved merger between AT&T and DirecTV is partially premised on the argument that the combined entity will be able to take advantage of scale economies to invest in a more expansive broadband network. If that argument is at least partially true, it signals to incumbent firms that anticompetitive abusive treatment of content providers is at best a short-term strategy with great long-term risks.

Perhaps the bigger issue in the competitive landscape for broadband is the role of Schumpeterian, or dynamic, competition. At any given moment, the market for high-speed broadband can be defined in a manner that seems to imply that most consumers are vulnerable to mistreatment by firms with monopoly power. Network quality improves through substantial capital investments, and short-run market power provides the primary incentive for such investments. At the same time, technology continually generates new types of content or methods of enhancing existing content, which in turn creates more demand for higher quality broadband networks. The recent history of the market appears to one of chicken-and-egg evolution with advances in networks supporting advances in content, and advances in content creating demand for more sophisticated networks.

Yet, if every time a substantial investment in network quality occurs, regulatory agencies treat the resulting picture of high concentration, when the market is defined by selecting the highest quality broadband service, as evidence of a lack of competition and therefore a justification for price regulation, then incentives to invest will be dulled. Indeed, the factions that demand price regulation will always be able to find an empirical justification given the economic relationship between quality, investment, and market concentration. This generates the paradoxical result that as the potential for significant network quality improvements increases—resulting in more choices for consumers among different levels of service quality—the threat of burdensome regulation based on a static vision of competition increases too.

5 Government Coercion

Another reason that the bridge analogy may seem incomplete is that it leaves out the problem of government coercion as a potential justification for the net neutrality norm. The government coercion problem provides net neutrality proponents with perhaps their strongest argument; but it is far from compelling in the end.

The threat of government coercion in the broadband service market arises in two settings: one is potentially desirable from a social welfare perspective, and the other is almost certainly undesirable. The potentially desirable form of coercion arises as

a response to internet piracy. Content owners have sued internet service providers, seeking a more active effort by the network providers to block pirate websites or sellers of knock-off products.²² Content owners invariably argue that the internet service providers have not done enough to ensure that copyrights or trademarks are not regularly infringed through their networks. In the case of trademarks, the networks typically enable the sale of cheap knock-offs of established brands that either directly infringe the established trademark, or exhibit sufficient similarity to have the same effect (trade dress infringement).

As a general matter, broadband networks should not be held strictly liable for copyright or trademark infringement. The networks provide a valuable social good by enhancing the flow of information in society. The positive spillover benefits that are provided by broadband networks probably outweigh the harms to copyright and trademark owners. Still, the remaining question is whether the owners of broadband networks have an obligation to monitor the theft of intellectual property and just how much effort such an obligation would entail on their parts.

The net neutrality norm enters here with an obvious application. Broadband network owners may prefer something like a network neutrality rule as a barrier to any claims of liability for copyright or trademark infringement. A broadband network could assert the neutrality norm as a justification for treating all content providers alike, and for reluctance to treat any of its content providers differentially. A network neutrality rule that is enacted into law might provide a powerful preemption defense to any broadband network that is sued for facilitating infringement of intellectual property.

In this setting, the net neutrality norm provides an overly broad principle that stands in the way of designing an optimal regulatory scheme. To monitor and control the infringement of intellectual property, broadband providers may have to adopt more aggressive measures than have been adopted to date. Courts that apply standard negligence principles are in an ideal position to determine the proper balance between risk and precaution on the part of broadband providers. A net neutrality regulation, which would be asserted in an effort to preempt a court from applying negligence principles, would effectively remove courts from an area of regulation in which they have traditionally operated, and in which they may be able to provide an appropriate solution to conflicting interests.

The other form of government coercion to which I referred earlier and which is not socially desirable, is state-directed control of internet content. The net neutrality norm arguably provides the benefit of enabling broadband networks to cite the norm (or regulation) as a justification for refusing to carry out the censorship aims of a particular government. Here, the norm appears to be unnecessary in a society in which the government is constrained to protect free speech, and would be wholly ineffective in a society in which the government is not so constrained.

In the United States, the First Amendment to the U.S. Constitution should form a sufficient legal basis for any broadband provider to resist efforts on the part of the

²² *MGM Studios, Inc. v. Grokster, Ltd.*, 545 U.S. 913 (2005); *BMG Rights Management (US) LLC et al. v. Cox Enterprises Inc. et al.*, case number 1:14-cv-01611, in the U.S. District Court for the Eastern District of Virginia.

U.S. government to control content—especially of a political nature—on the internet. It is not clear what a net neutrality norm could provide beyond the protection that is already provided by the constitutional guarantee of free speech. If anything, a net neutrality norm would likely be harmful in preventing the government from establishing, through either judicial or regulatory standards, optimal regulations to control abuses such as copyright infringement.

What if, instead of a governmental directive to the network to constrain speech, the network itself attempts to constrain speech? For example, suppose a network announces that it will impose a special charge on content providers who intend to transmit speech that the network regards as offensive to some particular person or group? The neutrality principle would arguably promote free expression in this special case by preventing the network firm from imposing discriminatory viewpoint-based charges.

However, there are several reasons to doubt the desirability of the neutrality norm even in this scenario. First, the neutrality norm goes too far by prohibiting all discriminatory charges rather than limiting its prohibition to viewpoint-based discriminatory charges. As in the scenarios that were considered earlier (anticompetitive abuse, preemption), the neutrality rule prohibits desirable forms of discrimination in an effort to root out potentially undesirable forms of discrimination. Second, competitive pressures (for customers, for capital) will constrain the incentives of network firms to discriminate on the basis of viewpoint. Indeed, if viewpoint discrimination were attractive to consumers, the heavy-handed internet censoring observed in China would not be met by massive efforts to circumvent it. Third, large broadband firms have valuable contracts with federal and state governments,²³ and would surely put these contracts at risk by engaging in viewpoint-based discrimination of content providers. These last two considerations suggest that the market will discourage broadband firms from engaging in viewpoint-based discrimination among content providers.

In a government that is not constrained (through both law and actual enforcement) to protect free speech, such as that of China, a net neutrality norm would be the feeblest of barriers against government coercion. The content providers themselves would face the threat of punishment if they were to violate the government's censorship rules. Broadband providers might or might not comply with a rule of net neutrality, but it would make no difference to the underlying problem of censorship by the government. Indeed, the net neutrality norm might facilitate the government's censorship by preventing broadband providers from taking any actions that might counteract the government's direct control over content providers. The net neutrality norm could then be paraded ostentatiously by the censoring government as a sign of its relative enlightenment.

²³ For example, on Comcast's government business, see <http://business.comcast.com/ethernet/industry-solutions/government>.

6 Conclusion

In theory, the implementation of the net neutrality norm might be socially desirable, and could be justified on the same grounds as the fair use doctrine in copyright law. However, in practice, the conditions that justify fair use in the copyright context do not appear to hold in the settings in which the net neutrality principle operates. Moreover, the cross-subsidization that is required by net neutrality generates a transfer from the relatively poor to the relatively rich. For every potential social gain that might be provided by the neutrality policy, an alternative, narrower policy exists that would be at least as effective and less likely to have harmful side effects. Efficiency and equity considerations provide no support for the net neutrality norm.

Acknowledgements I thank Michelle Connolly, Thomas Hazlett, David Sappington, Elliot Maxwell, Roger Noll, Scott Wallsten, Chris Yoo, and Lawrence White for very helpful comments. All errors and omissions are my own. I have also benefited from suggestions of participants at the Tech Policy Institute's Conference on Net Neutrality, January 8, 2016, in Washington, DC.

References

- Baumol, W. J., & Bradford, D. F. (1970). Optimal departures from marginal cost pricing. *The American Economic Review*, 60(3), 265–283.
- Becker, G. S., Carlton, D. W., & Sider, H. S. (2010). Net neutrality and consumer welfare. *Journal of Competition Law and Economics*, 6(3), 497–519.
- Clark, D. (2014). *Measuring Internet congestion: A preliminary report*. <https://ipp.mit.edu/sites/default/files/documents/Congestion-handout-final.pdf>.
- Economides, N., & Tåg, J. (2012). Net neutrality on the internet: A two-sided market analysis. *Information Economics and Policy*, 24(2), 91–104.
- Faulhaber, G. R. (1975). Cross-subsidization: Pricing in public enterprises. *The American Economic Review*, 65(5), 966–977.
- Gordon, W. J. (1982). Fair use as market failure: A structural and economic analysis of the Betamax case and its predecessors. *Columbia Law Review*, 82(8), 1600–1657.
- Hazlett, T. W., & Wright, J. D. (2012). The law and economics of network neutrality. *Indiana Law Review*, 45(3), 767–840.
- Hotelling, H. (1938). The general welfare in relation to problems of taxation and of railway and utility rates. *Econometrica*, 6(3), 242–269.
- Minasian, J. R. (1979). Indivisibility, decreasing cost, and excess capacity: The bridge. *Journal of Law and Economics*, 22(2), 385–397.
- Odlyzko, A. (1997). *A modest proposal for preventing internet congestion*. <http://www.dtc.umn.edu/~odlyzko/doc/modest.proposal.pdf>.
- Peitz, M., & Schuett, F. (2015). *Net neutrality and inflation of traffic*. TILEC Discussion Paper No. 2015-006. Available at SSRN: <http://ssrn.com/abstract=2573466> or <http://dx.doi.org/10.2139/ssrn.2573466>.
- Posner, R. A. (1971). Taxation by regulation. *The Bell Journal of Economics and Management Science*, 2(1), 22–50.
- Posner, R. A. (1981). *The economics of justice*. Cambridge, MA: Harvard University Press.
- Rawls, J. (1971). *Theory of justice*. Cambridge, MA: Belknap Press of Harvard University Press.
- Tirole, J. (1980). *The theory of industrial organization*. Cambridge, MA: MIT Press.
- Yoo, C. S. (2013). Is there a role for common carriage in an internet-based world? *Houston Law Review*, 51(2), 545–608.